

Geometric Filtration Using Proper Orthogonal Decomposition for Aerodynamic Design Optimization

David J. J. Toal,^{*} Neil W. Bressloff,[†] and Andy J. Keane[‡]
University of Southampton, Southampton, England SO17 1BJ, United Kingdom
and
Carren M. E. Holden
Airbus Operations Ltd., Bristol, England BS99 7AR, United Kingdom

DOI: 10.2514/1.41420

When carrying out design searches, traditional variable screening techniques can find it extremely difficult to distinguish between important and unimportant variables. This is particularly true when only a small number of simulations are combined with a parameterization that results in a large number of variables of seemingly equal importance. Here, the authors present a variable reduction technique that employs proper orthogonal decomposition to filter out undesirable or badly performing geometries from an optimization process. Unlike traditional screening techniques, the presented method operates at the geometric level instead of the variable level. The filtering process uses the designs that result from a geometry parameterization instead of the variables that control the parameterization. The method is shown to perform well in the optimization of a two-dimensional airfoil for the minimization of drag-to-lift ratio, producing designs better than those resulting from traditional kriging-based surrogate model optimization and with a significant reduction in surrogate tuning cost.

Nomenclature

C	=	fluctuation correlation matrix
C_D	=	drag coefficient
C_L	=	lift coefficient
F	=	matrix of snapshot fluctuations
l	=	design variable lower bounds
M	=	no. of snapshots
n	=	no. of sample points
p	=	hyperparameter determining smoothness
R	=	correlation matrix
S	=	snapshot ensemble
s	=	snapshot vector
u	=	design variable upper bounds
V	=	matrix of eigenvectors
x	=	design variable
y	=	objective function
α	=	modal coefficient
θ	=	hyperparameter determining correlation
Λ	=	vector of eigenvalues
λ	=	eigenvalue
μ	=	mean
σ	=	standard deviation
Φ	=	matrix of eigenfunctions

I. Introduction

THE optimization of complex geometries is prevalent throughout the field of aerodynamic design, from the optimization of two-dimensional airfoils to the optimization of complete wings and

aircraft. Designers are, however, restricted somewhat in the parameterization of such shapes. A typical airfoil parameterization, for example, may have anything from a few to more than 30 variables [1–5]. Parameterize an aircraft wing using a series of such airfoils and the number of variables can quickly be in the order of hundreds. Such large numbers of variables naturally restrict the ability of an optimization algorithm to achieve an optimal design given a limited budget for simulations.

To enable the efficient optimization of such geometries, it is therefore normal for a designer to employ some method of screening to identify those variables that influence the objective function most. With these variables identified, the optimization can proceed using the reduced variable set. Painchaud-Ouellet et al. [3], for example, reduced a 34-variable nonuniform rational B-spline (NURBS) parameterization of an airfoil to a total of 11 variables, and Song and Keane [6] reduced a 33-variable parameterization of an engine nacelle to seven important variables. Such large reductions in the complexity of an optimization are especially beneficial when the objective function is evaluated using an expensive high-fidelity computational simulation. The expense of such simulations, even with the growth of parallel computations, prohibits the exhaustive search of design spaces with large numbers of variables. In such problems, the reduction of the design space offered by variable screening is invaluable, allowing the designer to use a specified simulation budget more effectively.

Variable screening does, however, suffer from a number of significant problems. The cost of screening a complex optimization problem that requires evaluations using high-fidelity simulations may rival or exceed that of the optimization of the reduced variable set. Given that the screening procedure may not in fact return a definite set of variables that contribute most significantly to the change in objective function, the screening budget might be better used in the optimization of the original complete variable set. Variable screening also results in a reduction in the flexibility of the geometry parameterization, reducing the ability of the optimization to achieve certain high-performance designs that may only be found with the original, complete, variable set.

Optimization of complex geometries with large numbers of variables and a limited simulation budget therefore equates to a simple tradeoff. The designer can attempt to optimize the original parameterization and potentially achieve a high-performance design, though with a limited budget this optimal design will be difficult to locate. Alternatively, the designer can spend a proportion of their

Presented as Paper 6584 at the 26th AIAA Applied Aerodynamics Conference, Honolulu, HI, 18–21 August 2008; received 6 October 2008; revision received 21 January 2010; accepted for publication 23 January 2010. Copyright © 2010 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0001-1452/10 and \$10.00 in correspondence with the CCC.

^{*}Research Fellow, School of Engineering Sciences. Student Member AIAA.

[†]Senior Lecturer, School of Engineering Sciences.

[‡]Professor of Computational Engineering, School of Engineering Sciences.

simulation budget to screen out unimportant or insignificant variables and optimize the reduced design space using the remainder of the budget. The optimal design within the reduced space may be easier to locate, but the actual design may be significantly worse than that which could be obtained using the complete variable set. In the worst-case scenario, the variable screening procedure may indicate no significant difference in the importance of each variable. A suboptimal reduced set of variables may therefore be chosen, significantly hampering the optimizer's ability to achieve a good design.

An optimization strategy that combines the advantages of design space reduction associated with variable screening while retaining the majority of the flexibility of the original optimization would therefore be desirable. To this end, the authors have devised a strategy that combines surrogate modeling techniques and proper orthogonal decomposition (POD) in an attempt to optimize aerodynamic problems consisting of large numbers of variables but given only a limited simulation budget. Within the framework of this optimization strategy, proper orthogonal decomposition is used to reparameterize the problem in an attempt to filter out badly performing geometries.

Aerodynamicists already employ a form of geometric filtration, perhaps without even realizing it, for some aerodynamic optimization problems. Consider, for instance, the design of an airfoil for optimal performance at transonic speeds. Here, a designer uses their knowledge of existing airfoil shapes and bases the optimization on an existing supercritical airfoil, such as the RAE-2822. An optimization for a particular transonic flight condition therefore results in a series of perturbations to this baseline geometry, be it through the manipulation of NURBS [1] control points or through the addition of analytical functions [4]. The choice of this baseline airfoil has immediately filtered out badly performing designs and in doing so has reduced the design space before even commencing the parameterization.

Robinson and Keane [7] used this concept to construct an airfoil parameterization based on a series of orthogonal basis functions derived from an ensemble of supercritical airfoils. This parameterization technique operated on the principle that airfoils that perform well at a particular flight condition, the transonic regime in this case, had a number of common geometric features that could be extracted through the orthogonalization process and used for optimization. This is perfectly feasible in the case of 2-D airfoil design, where large databases of the performance of different airfoils exist, Abbott and Von Doenhoff [8] being a popular example. Problems arise when an optimization problem is encountered for which there is no such literature. There is, for example, no similar database of the performance of different wing-body fairings, the best designs of which could be used to construct a similar set of orthogonal bases. While the generation of such basis functions is undoubtedly a useful tool, this approach cannot be applied to every aerodynamics problem.

The following paper begins with an overview of the proposed geometric filtration procedure that is then applied to the optimization of a transonic airfoil for minimum drag-to-lift ratio. This test

optimization problem is investigated using three different optimization techniques. An extensive optimization of the airfoil is first performed directly using a genetic algorithm. The problem is then investigated using a traditional kriging-based approach, which is considered a direct competitor to the geometric filtration strategy. The paper then analyzes the performance of the geometric filtration strategy, considering the effect of the various optimization parameters on the quality of the final design obtained using the strategy. In an attempt to make the traditional kriging optimization strategy more competitive with respect to the overall cost of model tuning, the number of sample points used in the construction of the kriging model is restricted. The effect this restriction has on the optimization process and the speed of the optimization is then analyzed. Finally, the paper concludes by applying the same restriction to the geometric filtration strategy, resulting in a further reduction in surrogate model tuning time.

II. Optimization Via Geometric Filtration

A. Overview of the Complete Methodology

At its heart, the proposed optimization methodology consists of a basic surrogate modeling optimization strategy. This approach to optimization is particularly popular when objective function evaluations are extremely expensive. The surrogate model attempts to model the response of the objective function to changes in the design variables. A stochastic optimizer, for example, a genetic algorithm, can then draw on the surrogate model instead of the expensive computer simulation in an attempt to locate regions of optimal design. Typically, the expensive computer simulation is subsequently used to evaluate the true objective function at these positions. The true objective function can then be used to update the surrogate model, which can be searched again. This cycle is then repeated until some stopping criterion is reached.

The proposed optimization strategy shown in Fig. 1 also begins with an initial surrogate-based optimization: in this case, employing a kriging model to construct the surrogate. This initial optimization uses the designer's initial geometry parameterization and a proportion of the total simulation budget. This is followed by a reparameterization procedure using proper orthogonal decomposition, which attempts to both reduce the number of design variables and filter out badly performing designs. A secondary surrogate model optimization is then performed using this new geometry parameterization and the remainder of the total simulation budget.

In the context of the proposed optimization strategy, the initial surrogate model optimization followed by the POD performs a job similar to that of traditional variable screening, except that the screening is performed at the geometric level instead of the variable level. This prevents the reduction in variables being restricted to those of the original parameterization. It is hoped that the initial optimization provides a number of good designs that are then decomposed using POD into a series of orthogonal basis functions that capture the common features of these good designs. A similar process has been used by Kamali et al. [9] to reduce a design space but using points from a Latin hypercube sampling of the design space and not a subset of points from an initial optimization.

As shown throughout the literature [10,11], POD cannot accurately represent data outside of the initial snapshot ensemble. Here, this particular feature of POD combined with an initial search offers an advantage by considering only the current best geometries in the construction of the orthogonal bases; the bad geometries cannot be recreated through a combination of POD bases, and hence they have been filtered out. The second surrogate model optimization therefore has the benefit of a reduced number of variables with a minimal reduction in geometric flexibility.

B. Kriging

The implementation of the geometric filtration strategy considered within this paper uses a kriging-based optimization in both the initial and secondary optimizations. Although a different surrogate modeling method such as a simple polynomial or a radial

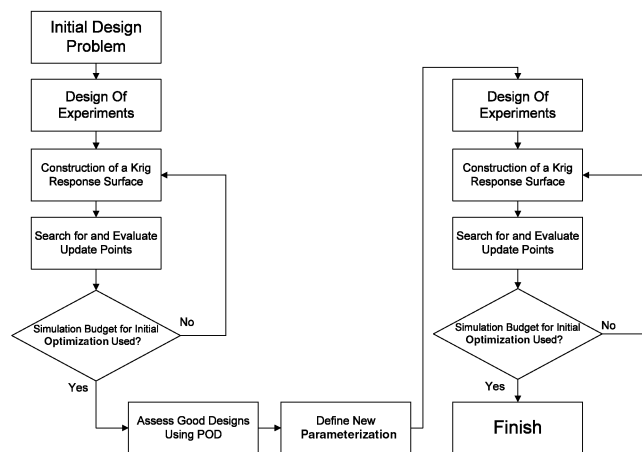


Fig. 1 Overview of the geometric filtration optimization methodology.

basis function could conceivably be used within this framework, kriging is used due to its ability to more accurately represent complicated responses while providing an error estimate of the predictor. First used by geologists in the estimation of mineral concentrations, it has since been popularized by Sacks et al. [12] in the creation of surrogate models of deterministic computational experiments and has been used successfully in the optimization of a number of different design problems [13–16].

To demonstrate the basic process of kriging, we consider the optimization of an objective function y , which is dependent on the vector of variables \mathbf{x} , which are the variables defining the initial geometry parameterization. The initial bound constrained optimization problem in d dimensions is therefore

$$\text{minimize } y(\mathbf{x}) \quad \text{subject to } \mathbf{l}_i \leq \mathbf{x}_i \leq \mathbf{u}_i \quad (1)$$

where \mathbf{l} and \mathbf{u} refer, respectively, to the lower and upper bounds of the variables \mathbf{x} . In general, the objective function values $y(\mathbf{x}_i)$ and $y(\mathbf{x}_j)$, which depend on d variables, will be similar if the distance between \mathbf{x}_i and \mathbf{x}_j is small. This can be modeled statistically by considering the correlation between two points as

$$\mathbf{R}_{ij} = \exp\left(-\sum_{l=1}^d \theta_l \|\mathbf{x}_i - \mathbf{x}_j\|^{p_l}\right) \quad (2)$$

where θ_l and p_l are known as the hyperparameters and determine, respectively, the rate of correlation decrease and the degree of smoothness in the l th direction. These hyperparameters, including a regression constant if required [17], are chosen to maximize the likelihood on the observed data set \mathbf{y} , where \mathbf{y} is a vector of n objective function values found by sampling the problem space.

The concentrated likelihood function [18],

$$-\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln(|\mathbf{R}|) \quad (3)$$

is evaluated by first calculating the mean,

$$\hat{\mu} = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}} \quad (4)$$

and then the variance,

$$\hat{\sigma}^2 = \frac{1}{n} (\mathbf{y} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}) \quad (5)$$

where $\mathbf{1}$ is an $n \times 1$ vector of ones and \mathbf{R} is the correlation matrix, the i th and j th elements of which are calculated using Eq. (2).

The concentrated likelihood is dependent only on the symmetric matrix \mathbf{R} and hence on the hyperparameters, which are then optimized to maximize the likelihood. Any optimization strategy could conceivably be used for this optimization problem, although a global optimizer is generally preferred due to the multimodal nature of the likelihood [19]. With the hyperparameters defined, the surrogate model can be used to predict regions that either minimize the model's prediction of the objective function or maximize its expected improvement [18].

C. Proper Orthogonal Decomposition

POD has been extensively used throughout the engineering community with regard to computational fluid dynamics. It has been used in the derivation of reduced-order models for the purposes of control [20–23] and optimization [24–26]. POD decomposes an ensemble of snapshots [27] of data into a set of optimal orthogonal basis functions of decreasing importance. The basis functions are optimal in the sense that no other basis functions will capture as much information in as few dimensions [28]. Applying POD to an ensemble of airfoils therefore results in a series of orthogonal basis functions similar to those of Robinson and Keane [7].

The decomposition process, as per Sirovich's [27] method of snapshots, begins with the definition of an ensemble of snapshot vectors \mathbf{S} . In the case of the geometric optimizations considered

within this paper, the snapshot ensemble is constructed from a series of M vectors consisting of the x and y coordinates of a number of designs selected from the initial optimization:

$$\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_M] \quad (6)$$

This matrix of snapshots is then decomposed into a mean $\bar{\mathbf{s}}$ and a matrix of the fluctuations of each snapshot from this mean:

$$\mathbf{S} = \bar{\mathbf{s}} + \mathbf{F} \quad (7)$$

where

$$\bar{\mathbf{s}} = \frac{1}{M} \sum_{i=1}^M \mathbf{s}_i \quad (8)$$

The orthogonal basis functions are calculated by considering the solution to the eigenvalue problem:

$$\mathbf{C}\mathbf{V} = \mathbf{\Lambda}\mathbf{V} \quad (9)$$

where $\mathbf{\Lambda}$ is a vector of eigenvalues, and the square symmetric correlation matrix \mathbf{C} is given by

$$\mathbf{C} = \mathbf{F}^T \mathbf{F} \quad (10)$$

The matrix of eigenvectors \mathbf{V} can then be used to calculate the matrix of eigenfunctions:

$$\mathbf{\Phi} = \mathbf{F}\mathbf{V} \quad (11)$$

where $\mathbf{\Phi}$ is a matrix of M eigenfunctions. These eigenfunctions, along with the corresponding vector of modal coefficients $\boldsymbol{\alpha}$, allow the fluctuations and hence the original snapshots to be reconstructed:

$$\mathbf{s}_i = \bar{\mathbf{s}} + \mathbf{\Phi} \boldsymbol{\alpha}_i \quad (12)$$

The advantage of POD is that not all of the M POD basis functions (eigenfunctions) are necessary to recreate the original snapshot ensemble to a required degree of accuracy. The cumulative percentage variation [29] can be used to define a reduced number of basis functions. The importance of each POD basis function is related to the relative magnitude of the corresponding eigenvalue; a large eigenvalue therefore indicates an important basis function. The cumulative percentage variation

$$\sum_{i=1}^N \lambda_i / \sum_{i=1}^M \lambda_i \times 100 \quad (13)$$

is therefore a measure of the combined importance of the first N basis functions. Using this simple calculation, a reduced number of basis vectors can be selected in order to meet a minimum required percentage variation. Using this reduced number of basis functions, the original snapshot vectors can be approximated by

$$\mathbf{s}_i \approx \bar{\mathbf{s}} + \sum_{n=1}^N \boldsymbol{\alpha}_{i_n} \boldsymbol{\phi}_n \quad (14)$$

where $\boldsymbol{\phi}_n$ and $\boldsymbol{\alpha}_{i_n}$ are the N most important basis vectors and corresponding modal coefficients. The modal coefficients corresponding to each of the original snapshot vectors can be calculated using the orthonormality property of the basis vectors.

The geometric filtration optimization methodology therefore moves from a surrogate modeling optimization based on the magnitude of the original design variables to one that considers the magnitude of the POD modal coefficients. The bounds of the secondary optimization are defined by the minimum and maximum modal coefficients of the original snapshot ensemble. The optimization problem therefore moves from the formulation of Eq. (1) in d dimensions to one in N dimensions:

$$\text{minimize } y(\boldsymbol{\alpha}) \quad \text{subject to } \boldsymbol{\alpha}_{\min} \leq \boldsymbol{\alpha}_i \leq \boldsymbol{\alpha}_{\max} \quad (15)$$

III. Optimization of a Transonic Airfoil

A. Description of the Test Problem

To demonstrate the effectiveness of the proposed strategy, the optimization of a 2-D airfoil is used as a test case. The RAE-2822 airfoil is parameterized using two NURBS curves, one each for the upper and lower surfaces, with the positions and weights of the controls points defining both NURBS curves optimized using a local Broyden–Fletcher–Goldfarb–Shanno (BFGS) optimization, as per Lépine et al. [1]. The resulting parameterization of the RAE-2822, shown in Fig. 2, consists of a total of 20 variables. The control points at the leading and trailing edges are fixed; the two control points on the line, $x = 0$, are only permitted to move vertically to maintain curvature continuity at the leading edge, whereas the remaining control points can move in both axes. Although the weightings of each control point remain fixed for the purposes of the following




Fig. 2 NURBS parameterization of the RAE-2822 airfoil.






Fig. 3 Comparison of a pressure distribution resulting from a VGK simulation to experimental data at Mach 0.725, Reynolds number of 6.5×10^6 and C_L of 0.658.



a)



b)

Fig. 4 Plots of a) pressure distribution and geometry of the original RAE-2822 airfoil with $C_D/C_L = 1.48 \times 10^{-2}$ or $C_L/C_D = 67.8$ and b) an example airfoil resulting from an optimization using a genetic algorithm followed by dynamic hill climber with $C_D/C_L = 1.08 \times 10^{-2}$ or $C_L/C_D = 93.0$.

optimizations, they could be permitted to vary, increasing the optimization to 31 variables.

The airfoil is optimized to minimize the drag-to-lift ratio at Mach 0.725, Reynolds number of 6×10^6 , and a fixed angle of attack of 2° using the viscous Garabedian and Korn (VGK) [30] solver. VGK comprises the generation of a 160×30 conformal mapping and a finite difference solution of the full potential equations coupled with integral methods for both laminar and turbulent boundary layers [31] (further details of the VGK method can be found in [30]). Figure 3 shows a comparison between a VGK simulation of the RAE-2822 airfoil and experimental data taken from Cook et al. [32] at Mach 0.725, a Reynolds number of 6.5×10^6 , and fixed lift coefficient of 0.658. VGK predicts the majority of the upper- and lower-surface pressures very well, with the exception of an underprediction of the shock strength and the exact shock position. The experimental data indicate a drag coefficient of 0.0107 at these flight conditions, whereas VGK predicts a drag coefficient of 0.0103: an error of approximately 4.2%. Pitching moment is also accurately predicted, with VGK predicting a coefficient of -0.0856 and the experimental data indicating a pitching moment of -0.090 : an error of 4.9%. Further comparisons of VGK results with experimental data can be found in the literature [30,33]. It should be noted at this point that all future plots of pressure distributions are to the same vertical scale as that of Fig. 3.

The speed of the VGK solver, approximately 1 s per simulation, offers a major advantage when analyzing the performance of the geometric filtration strategy. The speed of the solver allows extensive averaging to be carried out, thus giving a more accurate picture of the performance of both the traditional kriging and geometric filtration approaches to design optimization. At the selected design conditions of Mach 0.725 and a Reynolds number of 6×10^6 at 2° , the RAE-2822 is predicted to have a drag-to-lift ratio of 0.0148, equating to a lift-to-drag ratio of 67.8, and exhibits an upper-surface shock wave just before the midchord point (Fig. 4a).

B. Exhaustive Optimization Using a Genetic Algorithm

The speed of the objective function evaluations also allows a number of exhaustive searches of the design space using a genetic algorithm (GA). Stochastic methods such as the genetic algorithm provide a reliable way of locating the region of the global optimum within a design space; however, as they typically require a large number of function evaluations, it is not generally feasible to use such methods when expensive high-fidelity simulations are required. Applying a genetic algorithm to the above airfoil optimization problem provides a useful indication as to the true optimum that both the geometric filtration strategy and the traditional kriging strategy are attempting to attain.

The airfoil design problem was optimized using a GA followed by a dynamic hill climber (DHC), both implemented using the

OptionsMatlab [34] design exploration system. A budget of 10,000 function evaluations was used in the optimization, with the budget split evenly between the GA and DHC. The 5000 available function evaluations for the genetic algorithm equated to 100 generations of 50 points each. Although genetic algorithms are typically very good at locating the general region of the global optimum, they can be very slow to converge to a precise answer. The DHC is therefore used to converge the optimization toward a more accurate solution.

This extensive optimization was carried out a total of 10 times, producing an average C_D/C_L of 1.06×10^{-2} with a standard deviation of 4.59×10^{-4} , which equates to a C_L/C_D of 94.7: an improvement of some 28.4% over the original RAE-2822. Figure 4b provides an indication of the designs that this extensive optimization produced. Upon comparison with the original airfoil (Fig. 4b), one can clearly see a complete removal of the upper-surface shock wave, a reduction in the upper-surface pressure, and an increase in the lower-surface pressure, which results in an overall decrease in the drag-to-lift ratio. The optimization process has resulted in a reduction to the leading-edge curvature and an overall reduction in thickness. The trailing-edge camber of the airfoil has also been increased and the thickest region of the airfoil has moved forward, compared with that of the original RAE-2822 airfoil.

C. Standard Kriging-Based Optimization

The geometric filtration methodology considered within this paper comprises two kriging-based optimizations linked by a reparameterization of the design problem. It is therefore necessary to determine the performance of a standard kriging optimization with respect to the current airfoil design problem in order to provide a meaningful measure of the performance gains offered by the geometric filtration strategy.

Consider now a basic kriging optimization consisting of a total budget of 300 objective function evaluations, which is a value much more indicative of what would be available in a typical design optimization than the 10,000 evaluations used in the previous exhaustive search. Of this total simulation budget, one-third is used in the initial design of experiments (DOE), in accordance with the work of Söbester et al. [35], with the remaining budget reserved for updates to the surrogate model, which are evaluated in batches of 10. All updates to the kriging models used in each of the following investigations are based on each model's prediction of the objective function. A genetic algorithm is used to minimize the objective function predicted by the model, with the cluster centroids of the population selected as the update points.

The kriging model has its hyperparameters tuned after every other set of updates using a 5000-evaluation genetic algorithm followed by a dynamic hill climber. This strategy was found by Toal et al. [36] to

offer a significant reduction in tuning cost while having minimal impact on the performance of the optimization.

This traditional kriging-based optimization strategy is applied to the airfoil design problem a total of 50 times, with the random number seed used to generate the DOE changing each time. The search histories for each of these 50 optimizations are presented in Fig. 5a, with a typical airfoil resulting from one such optimization presented in Fig. 5b.

The 50 optimizations result in an average C_D/C_L of 1.16×10^{-2} with a standard deviation of 2.17×10^{-4} , which equates to a C_L/C_D of 86.4 and an improvement in the objective function of 21.7%. The improvement over the initial airfoil design is therefore not as significant as that obtained using the direct GA. However, one must note that the direct search used over 33 times the number of function evaluations. The kriging strategy should therefore be commended for attaining 76.3% of the improvement obtained by the exhaustive optimization with 3% of the simulation budget. Perhaps given enough time and updates, and using a better updating strategy such as expected improvement, one would expect the traditional kriging strategy to approach the global optimum [37,38] that the GA is assumed to have obtained.

Similar to the airfoils resulting from the exhaustive search (Fig. 4b), the kriging-based optimization results in a reduction in the upper-surface pressure and an increase in lower-surface pressure over the whole airfoil. The shock wave appears to have been removed, but there are a number of noticeable oscillations in the upper-surface pressure distribution, compared with that resulting from the direct genetic algorithm. This results in an overall decrease in drag-to-lift ratio, but one that is not quite as significant as that obtained using the genetic algorithm.

When one considers the average time taken for the traditional kriging optimization, approximately 22.0 h, one can observe a serious drawback associated with the use of this particular optimization strategy: the cost of hyperparameter tuning. Given the low cost of the simulations in this optimization, the majority of the time is due to the hyperparameter tuning process and, more specifically, the $\mathcal{O}(n^3)$ decomposition of the correlation matrix necessary in each evaluation of the likelihood. The size of this correlation matrix and hence the cost of every evaluation of the likelihood increases as more update points are added to the surrogate model. The cost of tuning the hyperparameters will therefore only increase as an optimization progresses and become a considerable bottleneck toward the end of the process.

A designer must therefore carefully consider the cost incurred by hyperparameter tuning relative to that of the objective function evaluations to determine if kriging is an appropriate optimization strategy for the problem. Taking the optimization of the RAE-2822 airfoil as an example, the cost of constructing the kriging model vastly outweighs the cost of each objective function evaluation, and a designer could use a simpler strategy but with a larger number of

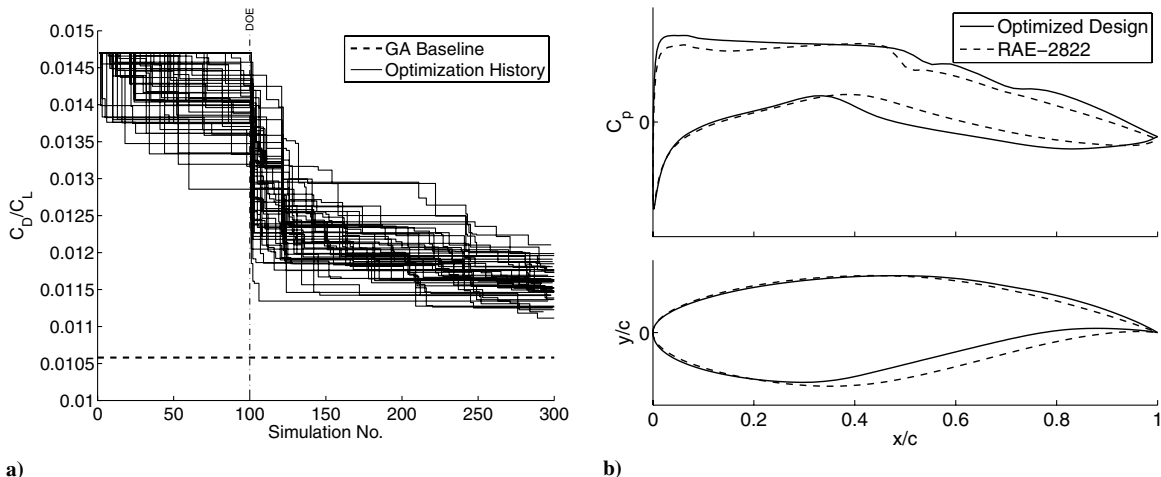


Fig. 5 Plots of a) optimization histories for each of the 50 traditional kriging airfoil optimizations and b) pressure distribution and geometry for an example airfoil resulting from a traditional kriging optimization process with $C_D/C_L = 1.16 \times 10^{-2}$ or $C_L/C_D = 86.2$.

evaluations, such as the direct GA. However, if the design problem is much more complex, for example, involving the optimization of a 3-D wing using Reynolds-averaged Navier–Stokes simulations, then the cost of each objective function evaluation may be in the region of tens of hours. For such a case, the extra time incurred through the hyperparameter tuning is more acceptable.

In terms of the total optimization time incurred, the implementation of a standard kriging strategy is therefore more practical given expensive objective function evaluations, but less so as the cost of the objective function decreases. Even so, the reduction in overall tuning cost, even for expensive simulations, is a worthwhile goal, allowing the designer to complete an optimization in a shorter time or allowing additional objective function evaluations within a similar time frame.

IV. Geometric Filtration

A. Initial Results

The optimization of the RAE-2822 airfoil using the geometric filtration strategy described in Sec. II, is now considered. In this initial investigation, a total budget of 300 simulations is again employed, with half of this budget used in the initial optimization.

In the initial kriging-based optimization, the available budget is split evenly between the initial DOE and the updates to the kriging model, instead of one-third for the DOE, as the initial response surface resulting from the use of a very small number of sample points is deemed inadequate. The secondary optimization, however, uses one-third of the available simulations in its DOE. A smaller ratio of DOE to updates is more reasonable for this optimization, as there are a smaller number of variables after the reparameterization, and hence fewer points are required to construct an adequate response surface.

In this particular implementation of the strategy, 30 airfoil geometries are selected from those generated during the initial kriging-based optimization. The airfoils are selected using a KMEANS [39] clustering algorithm in an attempt to maintain a measure of diversity between the selected airfoils. Upon decomposing the snapshot ensemble, the first N basis functions to capture more than 99.99% of the variation are selected as the POD modes for the secondary optimization. Such a high cumulative percentage variation ensures that the original airfoils contained within the snapshot ensemble can be recreated to a high degree of accuracy. This high accuracy means that the original 30 airfoils exist within the reparameterized design space, thereby allowing what is essentially 30 additional “free” design points to be added to the design of experiments of the secondary optimization.

In summary, the optimization begins with a DOE of 80 points, followed by 70 updates to the kriging model in batches of 10. From

this initial optimization, 30 geometries are selected and reparameterized using POD. The secondary optimization then optimizes the modal coefficients, beginning with a DOE of 50 points (plus the 30 airfoils from the snapshot ensemble) and followed by 100 updates to the model, again in batches of 10. The hyperparameters are once again tuned after alternate updates.

Once again, 50 different optimizations are carried out, with the random number seed used in the Latin hypercube changing each time. The histories for each of these optimizations are presented in Fig. 6a. The 50 optimizations result in an average drag-to-lift ratio of 1.13×10^{-2} with a standard deviation of 3.76×10^{-4} , equating to an average lift-to-drag ratio of 89.0. Applying the geometric filtration optimization strategy to this design problem has therefore resulted in an average of 84.1% of the improvement obtained with the genetic algorithm. From Fig. 6a it is also observed that a number of the final designs actually achieve a drag-to-lift ratio similar to that of the direct GA. The average results of these optimizations along with those of traditional kriging and the GA are presented together in Table 1.

Figure 6b shows a typical airfoil resulting from an optimization. Upon comparison with the airfoil in Fig. 5b, resulting from the traditional kriging optimization, one can observe the removal of the shock wave but with a reduction to the oscillations in upper-surface pressure. The reduction in the upper-surface pressure and increase in the lower-surface pressure has been maintained. The optimization has also resulted in an airfoil with a more significant reduction to the leading-edge curvature, something that was observed with the direct GA (Fig. 4b) but that is not as pronounced in Fig. 5b.

Upon comparison of the optimization histories of Figs. 5a and 6a, the application of geometric filtration appears to result in an increase in the variance of the final designs. The standard deviation of the drag-to-lift ratio for the traditional kriging strategy was 2.17×10^{-4} , whereas geometric filtration resulted in an increased standard deviation of 3.76×10^{-4} . By employing a subset of the good designs resulting from an initial kriging-based optimization, geometric filtration attempts to filter out undesirable features. In the process, the ability to recreate some good designs in the original space may also be lost. The quality of the snapshot designs is therefore important to the geometric filtration process: the more consistent the snapshots, the more consistent the final designs.

Figures 7a and 7b show the first six basis functions employed in the optimization that resulted in the airfoil of Fig. 6b. The first two of these functions govern changes to the aft of the upper surface and the majority of the lower surface. The fourth, fifth, and sixth basis functions again govern changes to the lower surface, particularly close to the leading edge. This corresponds well to the results obtained with the exhaustive search (Fig. 4b), which demonstrates that changes to the lower surface were significant in the reduction in the drag-to-lift ratio. Particularly, the exhaustive search made sig-

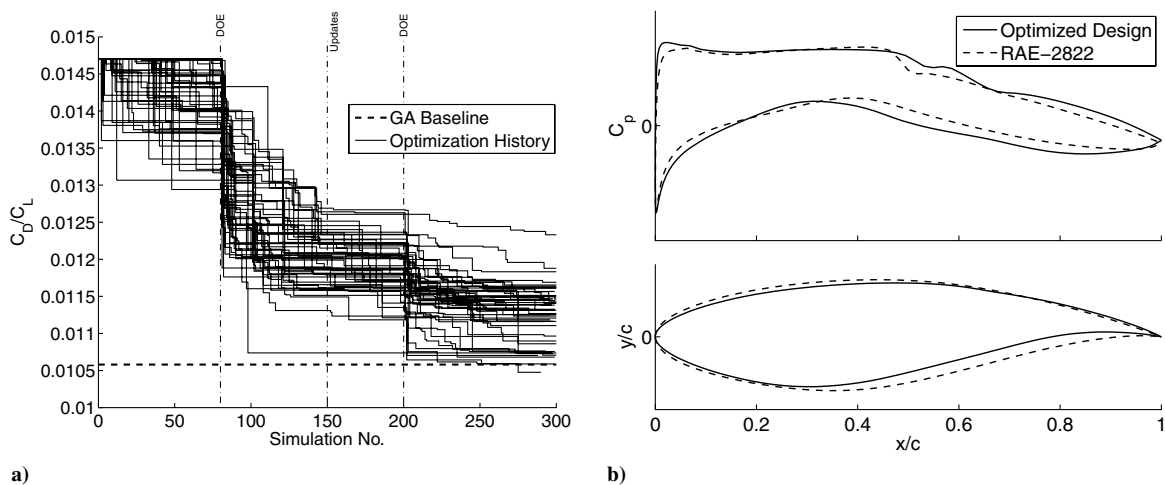


Fig. 6 Plots of a) optimization histories for each of the 50 optimizations using the geometric filtration strategy with an equal simulation budget for both the initial and secondary optimizations and b) pressure distribution and geometry for an example airfoil resulting from the geometric filtration optimization process with $C_D/C_L = 1.13 \times 10^{-2}$ or $C_D/C_L = 88.9$.

Table 1 Comparison of the three optimization strategies with respect to the optimization of the RAE-2822 for minimum drag-to-lift ratio

Strategy	No. of evaluations	Mean C_D/C_L	Std C_D/C_L	Mean C_L/C_D	Std C_L/C_D	Time, h
Genetic algorithm	10,000	1.06×10^{-2}	4.59×10^{-4}	94.7	4.60	2.78
Traditional kriging	300	1.16×10^{-2}	2.17×10^{-4}	86.4	1.62	22.0
Geometric filtration	300	1.13×10^{-2}	3.76×10^{-4}	89.0	2.98	4.15

nificant changes to the lower-surface leading edge, which is something that did not occur with the kriging strategy (Fig. 5b). However, the basis functions show that the filtration process has identified that changes to the leading edge are important. The basis functions also demonstrate that changes to the upper surface are less important, which is a result that is again borne out by the results of the exhaustive search. Each of the basis functions appears to contribute to the final design. Even though the higher modes cause much smaller changes to the geometry, such change can be significant when optimizing wetted surfaces.

The application of geometric filtration also results in a significant reduction in the total optimization time, which drops from an average of 22.0 h for a traditional kriging-based optimization to 4.15 h. As the hyperparameter tuning process is again responsible for the majority of the optimization time, geometric filtration has resulted in a substantial reduction in tuning cost. This reduction in cost can be explained when one considers the actual tuning process in each optimization, particularly the size of the correlation matrix \mathbf{R} , used to calculate the likelihood equation (3).

In the traditional kriging-based optimization, the correlation matrix begins 100×100 in size and steadily grows as more update points are evaluated, until the matrix is 280×280 in size. Geometric filtration, however, begins with a correlation matrix of 80×80 , which grows to 140×140 by the final tune of the initial optimization. Following the reparameterization process, 110 of these design points are filtered out and 50 new points are added from the design of experiments of the secondary optimization. The correlation matrix therefore grows steadily from 80×80 to 160×160 during the final hyperparameter optimization. When geometric filtration is employed, the cost of the $\mathcal{O}(n^3)$ factorization in each likelihood evaluation is reduced, which reduces the overall cost of the hyperparameter optimization.

In conclusion, the optimization of the RAE-2822 using this implementation of the geometric filtration strategy has resulted in better designs than the traditional kriging strategy for a considerably reduced tuning overhead but with a slight reduction in the consistency of these designs. In the following sections, the authors attempt to examine the implications of each of the different parameters of geometric filtration on the strategy's overall performance. The implications of the size of the initial and secondary optimization budgets, the size of the POD snapshot ensemble, and the number of POD basis functions selected are considered.

B. Implications of Initial and Secondary Optimization Budget

Of the many parameters influencing the setup of the geometric filtration strategy, the ratio of the total optimization budget used in the initial and secondary optimizations could be considered one of the most significant. The designs evaluated during the course of the initial optimization are used in the reparameterization process; hence, one would expect that the more effort expended in this initial optimization, the better the snapshot airfoils and the more important the features captured by the resulting basis functions. As always, optimization is a battle between exploration and exploitation; in this case, increasing the effort applied to exploration in the initial optimization reduces the budget available for the secondary optimization to exploit the reduced set of variables resulting from the reparameterization. To investigate this we now consider two optimizations employing different simulation budgets in the initial kriging-based optimization.

The first uses one-third of the total simulation budget in the initial optimization and the second uses two-thirds. The other parameters remain identical to that of the initial geometric filtration optimization. The DOE of both of the initial optimizations use approximately half of the total budget for that optimization, and the secondary optimization uses one-third. The alternate tuning strategy is again employed with 30 airfoils selected to form the snapshot ensemble. Once again, the first N POD basis vectors were selected to reproduce 99.99% of the cumulative percentage variation. The results for each of these optimizations along with those for the initial optimization are presented in Table 2. It should be noted that the additional optimization using two-thirds of the simulation budget in the initial optimization begins with a DOE identical to that of the traditional kriging strategy. This particular case therefore provides a direct performance comparison with the traditional strategy.

The results presented in Table 2 indicate that although geometric filtration consistently outperforms the traditional kriging approach, in terms of the average final objective function, the size of the initial optimization budget has an effect on both the quality and consistency in the final designs. Reducing the initial budget (run A) increases both the average objective function and the variation in the designs. Increasing the ratio to one-half (run B) results in better designs on average and a reduction in the variation. A further increase to two-thirds (run C) results in a slight increase in the average final objective function and a further reduction in the variation. The results of this optimization serve to reinforce the performance improvement offered by geometric filtration over the traditional kriging strategy.

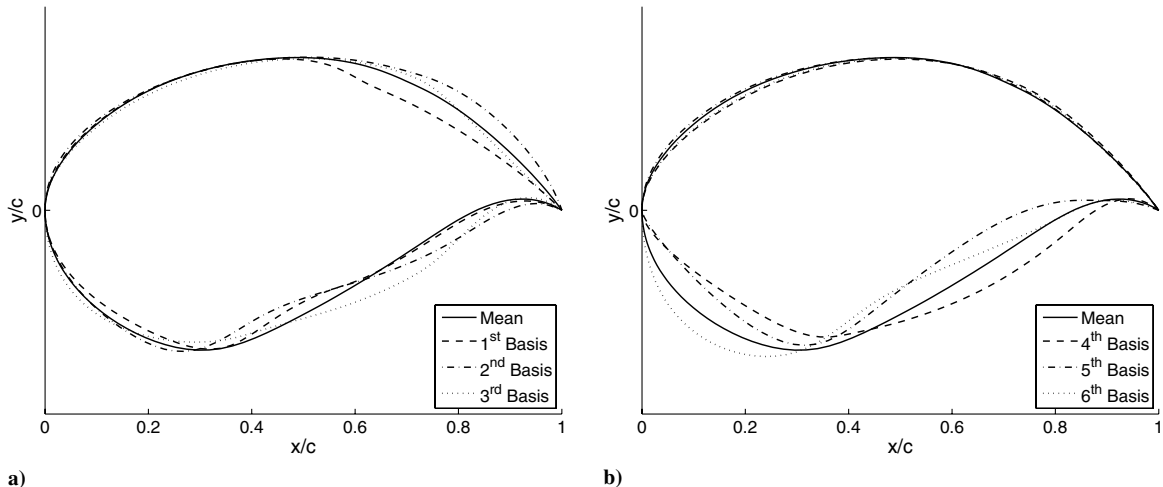
**Fig. 7 Examples of the six most important basis functions and their effect on the mean airfoil shape.**

Table 2 Comparison of geometric filtration optimizations using different ratios of simulation budgets in the initial and secondary optimizations

Run	Initial: secondary budget	Mean C_D/C_L	Std C_D/C_L	Mean C_L/C_D	Std C_L/C_D	Time, h
A	1:2	1.14×10^{-2}	4.14×10^{-4}	88.0	3.15	5.50
B	1:1	1.13×10^{-2}	3.76×10^{-4}	89.0	2.98	4.15
C	2:1	1.14×10^{-2}	2.69×10^{-4}	88.0	2.08	5.11
D	2:1 ^a	1.13×10^{-2}	2.65×10^{-4}	88.6	2.06	4.95

^aSecondary optimization containing no additional DOE points, only updates to a surrogate model defined using the snapshot ensemble points.

Both run C and the traditional kriging optimization begin from an identical DOE, but implementing the reparameterization of the problem results in better designs, on average.

From the optimization histories of Figs. 6a, 8a, and 8b one can observe that as more effort is applied to the initial optimization, the quality of the designs obtained by the end of this optimization improves. As the snapshot ensemble contains better designs, the resulting filtration is able to retain better geometric features and produce better designs in the secondary optimization. It can also be observed that the variation in the final designs is related to the variation in the designs obtained from the initial optimization. A large spread in the quality of the initial optimization produces a corresponding spread in the final design. As Table 2 demonstrates, the more effort that is applied to the initial optimization in order to obtain a better set of snapshots, the more the consistency between the final designs approaches that of the traditional kriging strategy.

Increasing the initial optimization reduces the ability to effectively optimize the reparameterized problem. This results in the observed increase in the average objective function when two-thirds of the total simulation budget is used in the initial optimization. This can be countered by reducing the size of the DOE evaluated at the beginning of the secondary optimization. As the basis functions have been chosen to represent 99.99% of the cumulative percentage variation, these points effectively exist within the design space of the secondary optimization. Using only these points as the DOE of the secondary optimization frees up more simulations to exploit the kriging model. As seen by the final result of Table 2, this produces designs with an average objective function close to that of the initial geometric filtration optimization, but with a much smaller variation in the final designs.

C. Implications of Snapshot Ensemble Size

We next consider the impact of altering the size of the snapshot ensemble on the performance of the geometric filtration optimization strategy. The effect of this parameter can be observed by analyzing the results of a number of additional optimizations, varying the ensemble size in each. For each of these additional optimizations the

same initial optimization using the original airfoil parameterization is used, ensuring that the impact of ensemble size is completely apparent. Snapshot ensembles of 10, 20, 30, and 50 are all considered, the results of which are presented in Table 3.

Increasing the size of the snapshot ensemble, especially if a clustering algorithm is used to select the design points, increases the diversity of ensemble designs, hence increasing the number of basis functions required to represent 99.99% of the cumulative percentage variation. The number of basis vectors required increases with the number of snapshot airfoils; however, after about 30 snapshots the increase is not as pronounced. Even though further snapshots are adding some information, they are similar to those already included in the ensemble, and so fewer additional basis vectors are required in order to fully represent them.

Increasing ensemble size decreases the average objective function of the final design. As suggested by the increase in the number of bases, adding more designs to the ensemble increases the variation of geometries that can be represented, hence increasing the size of the design space. A larger, more versatile, design space results in the capacity to create better geometries and thus a greater reduction in objective function.

In Fig. 9a one can observe the progress of the secondary optimization based on the POD basis vectors defined using an ensemble of 10 snapshot airfoils. Using the reparameterization resulting from these snapshots in the secondary optimization results in less improvement in objective function compared with that of Fig. 9b. Here, the reparameterization is based on 50 snapshots, resulting in a more flexible reparameterization, and therefore produces a greater improvement in objective function.

One should note that as the number of basis functions increases, so does the number of variables in the secondary optimization, therefore increasing the difficulty of this optimization given a limited simulation budget. However, as a large cumulative percentage variation is used, the ensemble design points can be added to the initial DOE of the secondary optimization. A larger ensemble size therefore means a larger DOE in the secondary optimization, and the negative impact of the increase in dimensionality is somewhat mitigated. The larger number of points in the secondary optimization does, however, result

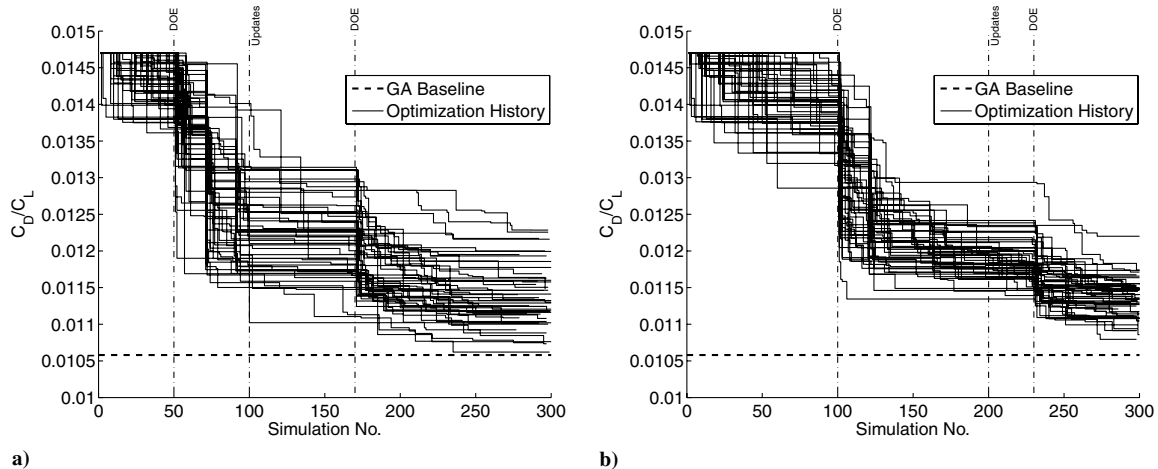


Fig. 8 Optimization histories for each of the 50 optimizations using the geometric filtration strategy with a) one-third and b) two-thirds of the total simulation budget used in the initial optimization.

Table 3 Results of geometric filtration optimizations with increasing snapshot ensemble size

Size of snapshot ensemble	Mean C_D/C_L	Std C_D/C_L	Mean C_L/C_D	Std C_L/C_D	Time, h	Mean no. of basis functions
10	1.16×10^{-2}	3.36×10^{-4}	86.4	2.25	2.85	8.5
20	1.14×10^{-2}	3.46×10^{-4}	87.8	2.65	3.64	12.7
30	1.13×10^{-2}	3.76×10^{-4}	89.0	2.98	4.15	14.1
40	1.13×10^{-2}	3.48×10^{-4}	88.8	2.71	4.52	14.8
50	1.11×10^{-2}	3.48×10^{-4}	90.1	2.82	5.87	15.4

in an increase in hyperparameter tuning time due to the larger correlation matrix, as shown by the results of Table 3.

D. Implications of POD Basis Selection

In the previous analysis it has been assumed that a cumulative percentage variation of 99.99% is enough to reproduce the snapshot airfoils to such a degree that they exist within the design space of the secondary optimization. We now consider the relaxation of this limit and the ability of the remaining POD basis functions to recreate the geometry of the original snapshot airfoils and, more important, the effect of any error in the geometric reconstruction on the correct prediction of lift and drag.

We begin by analyzing the recreation of 50 snapshot airfoils from a typical ensemble used in the previous analysis. Figure 10a shows the change in the average error in the recreation of the geometry of the original snapshots as the number of basis used in the recreation increases. With only a few basis vectors the original airfoils cannot be

accurately recreated, but increasing the number of basis functions rapidly reduces this error. For this particular ensemble, the first 16 POD basis functions result in a cumulative percentage variation of greater than 99.99%, which equates to a small average rms reconstruction error of 3.64×10^{-5} .

Let us now consider the effect of this error on the aerodynamic properties used in the current airfoil optimization. Figure 10b shows the average absolute error in both C_L and C_D as the number of basis vectors used to recreate the original ensemble of airfoils is increased. As expected, an increase in the accuracy of the geometry corresponds to a reduction in the error of both lift and drag. However, this reduction is not as rapid as that of the reduction in geometric error. It is generally recognized that even small changes to a geometry can significantly alter these aerodynamic coefficients, especially in the transonic flow regime, where slight changes to geometry can effect shock strength and position and hence lift and drag. A relatively small error in the recreation is therefore required to produce accurate aerodynamic responses.

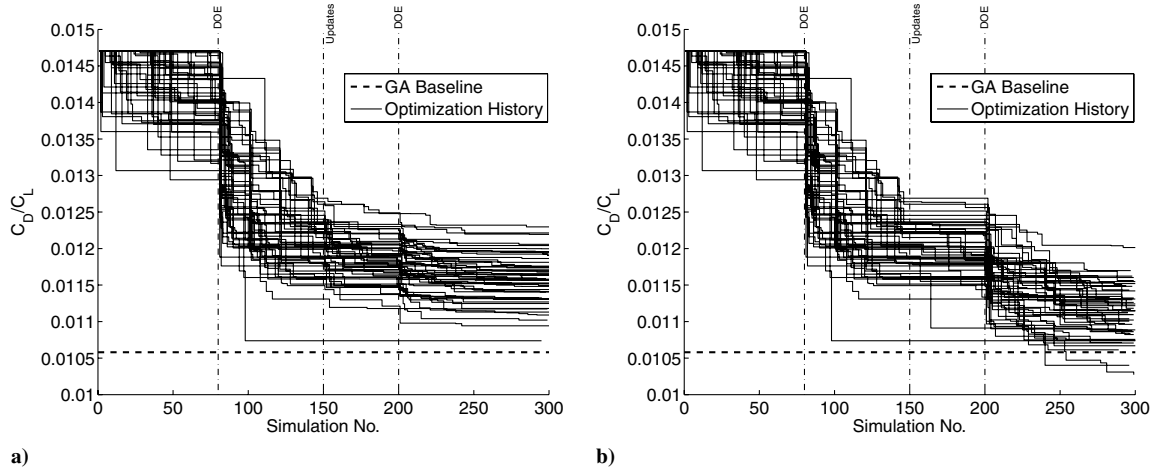


Fig. 9 Optimization histories for the geometric filtration strategy with the reparameterization based on a) 10 and b) 50 airfoil snapshots.

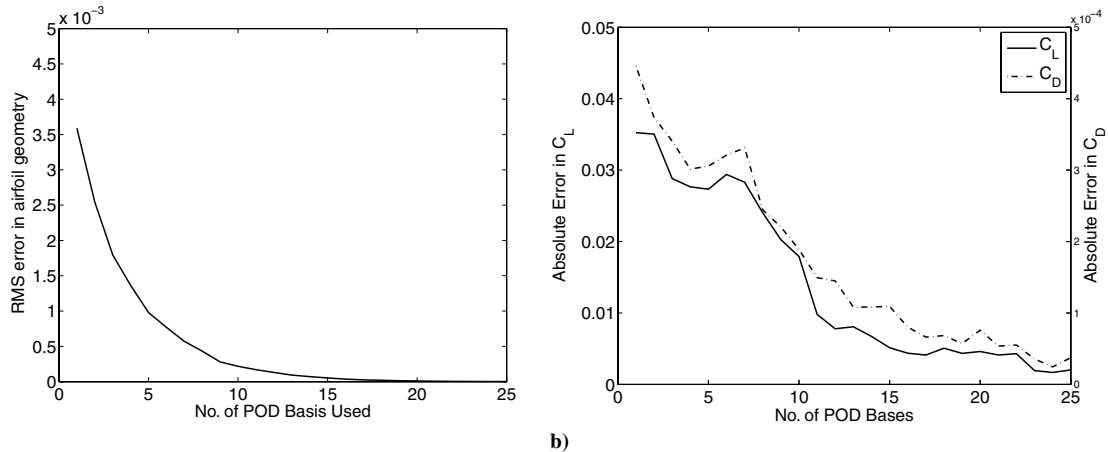


Fig. 10 Plots of a) demonstration of the accurate recreation of the original snapshot airfoils with increasing number of POD basis functions and b) the effect this has on the error in lift and drag coefficients.

Table 4 Performance results for the standard kriging optimization approach with varying sample size

Kriging sample size	Mean C_D/C_L	Std C_D/C_L	Mean C_L/C_D	Std C_L/C_D	Time, h
50	1.21×10^{-2}	6.25×10^{-4}	82.6	4.22	1.49
75	1.15×10^{-2}	4.15×10^{-4}	86.8	3.01	2.67
100	1.14×10^{-2}	2.03×10^{-4}	88.0	1.56	4.31
150	1.13×10^{-2}	1.59×10^{-4}	88.2	1.24	8.07
200	1.15×10^{-2}	1.93×10^{-4}	87.3	1.45	11.59
300	1.16×10^{-2}	2.17×10^{-4}	86.4	1.62	21.97

The results presented in Fig. 10b show that when 16 basis vectors are used, the errors in C_L and C_D are, respectively, 4.40×10^{-3} and 8.02×10^{-5} , on average. An average C_D error of just under one drag count could be considered acceptable for the purposes of including the snapshot objective function values within the DOE of the secondary optimization. However, including these objective function values in a secondary design that has used a smaller number of basis vectors would be unwise unless provision is made for the possible error in these results. Without such a provision the response surface would not accurately represent the true response of the objective function to changes in the modal coefficients, therefore introducing discontinuities in the response surface that may hamper the secondary optimization.

V. Performance of a Restricted Kriging Data Set

Throughout the previous analysis the impact of the reduction in the number of sample points on the total optimization time has been very evident. Reducing the number of sample points reduces the cost of evaluating the concentrated likelihood, thereby reducing the total optimization cost from 22.0 h using the traditional kriging strategy to an average of 4.2 h for the initial geometric filtration optimization.

Reconsidering the traditional kriging strategy, it is obvious that maintaining a constant restricted number of sample points will reduce the overall tuning time. However, what effect does such a restriction have on the strategy's overall performance and how does the performance compare with that of geometric filtration at an equivalent total optimization cost? To this end, five additional optimizations are performed using the traditional kriging optimization strategy but restricting the number of sample points within the design space to 50, 75, 100, 150, and 200 points. To do this, only the best points are retained as the optimization progresses. For example, restricting the design space to 150 points results in the optimization progressing normally after the DOE until the limit of 150 points is reached. Once this limit is reached, the best 150 points (in terms of their objective function) are retained. The results of this investigation are presented in Table 4 and graphically in Figs. 11a and 11b.

Figure 11a demonstrates the large reduction in the total optimization time achievable by restricting the maximum number of points within the response-surface model. It can be observed that restricting the number of points to a maximum of 100 results in a total optimization time that is much more competitive with that of the initial geometric filtration results. However, geometric filtration still performs better, achieving an average final objective function of 1.13×10^{-2} , whereas the restricted kriging strategy achieves 1.14×10^{-2} .

The restriction in the number of sample points also results in an improvement in the overall performance of the traditional kriging strategy. The average final objective function decreased from 1.16×10^{-2} with the traditional strategy to 1.13×10^{-2} when a restriction of 150 points is used. This is similar to that obtained by the initial implementation of the geometric filtration strategy, but for almost twice the tuning cost.

VI. Application of a Restricted Data Set to Geometric Filtration

As the geometric filtration optimization strategy is based upon a series of surrogate model optimizations, any techniques that can be applied to a normal surrogate model can also be applied to geometric filtration. The initial or secondary optimizations could make use of gradient-enhanced surrogates, alternative updating formulations or even variable-fidelity simulations through cokriging.

We have observed that by restricting the number of sample points within a design space of a traditional kriging strategy, the optimization time can be reduced and performance can be increased. In Sec. IV, increasing the initial optimization simulation budget was found to increase the consistency in the designs obtained, but at the cost of a slight reduction in the average objective function of the final design. This was countered by using only the points from the snapshot ensemble as the basis for the second kriging response surface. Increasing the size of the snapshot ensemble was observed to increase the variation in the geometries used, hence increasing the flexibility of the parameterization used in the secondary




Fig. 11 Plots of a) total wall time for a kriging optimization as the number of sample points used is reduced and b) average and standard deviation of the objective function of the final design as the number of sample points used is reduced.

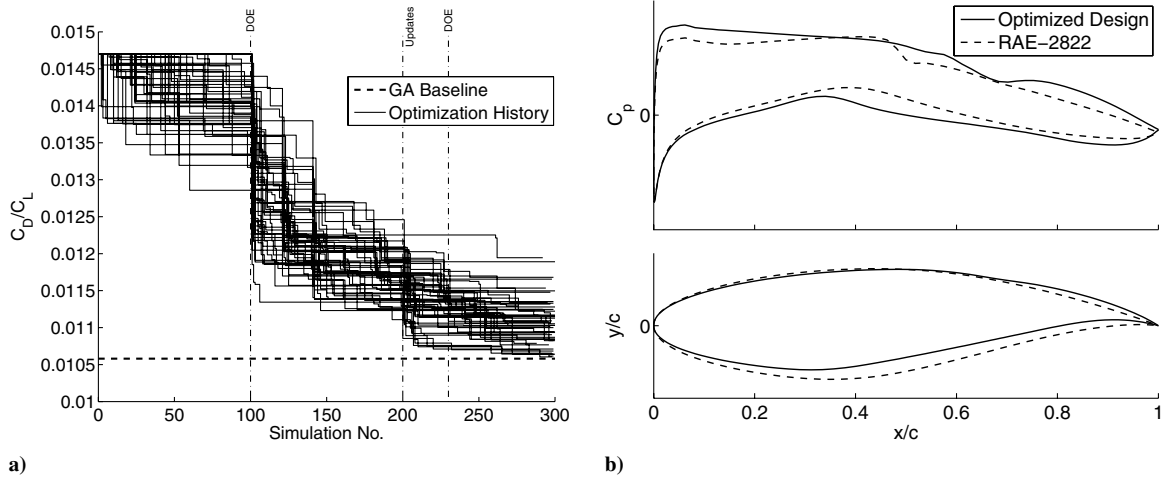


Fig. 12 Plots of a) optimization histories for each of the 50 geometric filtration optimizations using a restricted number of sample points and b) pressure distribution and geometry for an example airfoil with $C_D/C_L = 1.11 \times 10^{-2}$ or $C_L/C_D = 90.3$.

optimization. However, this came at the cost of an increase in the number of variables in the secondary optimization and a slight increase in overall optimization time.

We now consider an implementation of the geometric filtration strategy that takes all of the observations made into account. Once again, 300 simulations are used in the optimization. To attempt to increase the consistency between the final designs, two-thirds of the total simulation budget is used in the initial optimization, with half of these reserved for the DOE. The snapshot ensemble comprises a total of 50 airfoils in an attempt to increase the flexibility of the parameterization used in the secondary optimization. To improve the performance of the secondary optimization, the additional DOE is once again neglected and the remaining 100 simulations are used as updates to a kriging model initially defined using only the snapshot points. Once again, N basis vectors are selected to represent at least 99.99% of the variation.

The above implementation of the geometric filtration strategy obtained an average C_D/C_L of 1.11×10^{-2} , which equates to a C_L/C_D of 90.3. By adjusting the setup of the strategy in this way, 88.6% of the improvement in objective function obtained by the exhaustive search has now been obtained with approximately 3% of the simulation budget. Using a larger proportion of the total simulation budget in the initial optimization results in an increase in the consistency between the final designs, with a standard deviation of the final objective function reducing to 3.20×10^{-4} from 3.76×10^{-4} . This is still not quite as good as that obtained during the original investigation into the impact of initial optimization size and is probably due to the negative impact of the increase in the number of variables in the secondary optimization.

The optimization histories of Fig. 12a graphically demonstrate the clear improvement over both the traditional kriging strategy and the previous implementations of geometric filtration. Figure 12b presents an example airfoil resulting from a typical optimization. The upper-surface shock wave has again been removed, but there is a further reduction in upper-surface pressure and an increase in lower-surface pressure, resulting in a smaller drag-to-lift ratio.

Not only has this particular implementation of the strategy resulted in consistently good designs, but the restriction to the number of sample points has resulted in a substantial reduction in overall optimization time, with each optimization taking an average of 2.51 h. The total optimization time has in fact been reduced to such a degree by the implementation of geometric filtration that it is now faster than the extensive optimization using the GA and DHC, which took approximately 2.78 h to complete.

Applying the lessons learned from the analysis of the performance of the geometric filtration strategy has therefore resulted in an optimization strategy that not only comprehensively outperforms a traditional kriging strategy but does so for a considerable reduction in overall optimization time, even when applied in conjunction with a solver as fast as V GK.

VII. Conclusions

An optimization strategy involving two kriging-based response-surface optimizations and a POD based reparameterization has been introduced and applied to the optimization of an airfoil for minimum drag-to-lift ratio. This strategy, termed geometric filtration, was found to consistently outperform a traditional kriging-based optimization, producing better designs for a considerable reduction in overall optimization cost.

The geometric filtration strategy applies an initial kriging response-surface model optimization to the original problem. From the results of this optimization, a number of good design points are selected to form a snapshot ensemble for the purposes of proper orthogonal decomposition (POD). The POD basis functions then act as a reparameterization of the original problem, filtering out badly performing designs and reducing the number of variables. A secondary kriging response-surface based optimization is then carried out in which the magnitudes of the modal coefficients of a series of POD basis functions are optimized.

An in-depth analysis of the impact of the size of the initial optimization budget, the size of the POD snapshot ensemble, and the number of basis functions selected on the strategy's performance has been carried out. This indicated that increasing the size of the initial optimization budget effects the consistency of the final designs obtained. Increasing the number of snapshot airfoils was found to produce better final designs by increasing the flexibility of the secondary optimization parameterization. Selecting a reduced number of basis functions was found to reduce the cost and complexity of the secondary optimization, but at the expense of a reduction in flexibility of the parameterization. It was also observed that the number of basis functions used in the secondary optimization should be capable of capturing a large percentage of the cumulative variation if the objective function values from the snapshot airfoils are to be added to the DOE of the secondary optimization. Without a sufficient number of basis functions, the inclusion of these points may introduce discontinuities into the response surface, hampering the optimization.

A restriction on the number of sample points was introduced and applied to the traditional kriging optimization process to make it more competitive in terms of the total optimization time. This was found to not only reduce the total cost of each optimization, but to improve the average objective function of the final design in some cases.

The flexibility of geometric filtration with regard to the implementation of performance enhancing features used in traditional surrogate models was demonstrated with the application of a restriction in the number of sample points. The previous observations made regarding the effect of an increased initial optimization budget and an increase in the size of the snapshot ensemble were also applied. The resulting optimizations achieved a further improvement in the average objective function of the final designs, attaining 88.6% of the improvement obtained by a direct

genetic algorithm for 3% of the total simulation budget and for 11.4% of the total optimization time of the traditional kriging strategy.

Despite the investigations carried out within the current paper, there remain a number of areas for future work. The sensitivity of the geometric filtration strategy to the method used to select the snapshot airfoils remains unclear. Could a simpler strategy whereby the best M designs are selected prove just as effective, or is the diversity maintained by the KMEANS clustering algorithm a necessity? The performance of a recursive implementation of the presented strategy also requires investigation. The successive reduction in the number of variables is attractive, but would any increase in the variability in the quality of the resulting designs prove detrimental to performance? The modality of the secondary design space has not been considered; could a gradient decent technique prove just as successful or is a global optimization still required? Finally, no constraints were considered in the presented airfoil design problem, and the methodology for selecting potential snapshots given the presence of constraints requires investigation.

The choice of optimization strategy generally depends on the cost of evaluating the objective function. As engineering design problems often involve expensive simulations, direct optimization becomes impractical and the use of surrogate models becomes a necessity, but the associated model construction cost can represent a considerable bottleneck. The reduction in tuning cost demonstrated by the geometric filtration strategy significantly reduces this bottleneck, allowing for a faster overall optimization or more actual function evaluations in a similar time, while producing better final designs and a reduction in the number of variables.

Acknowledgments

The presented work was undertaken as part of an Airbus-funded activity. The authors would like to thank A. Forrester, A. Söbester, and I. Voutchkov of the University of Southampton for their advice and input.

References

- [1] Lépine, J., Guibault, F., and Trépanier, J., "Optimized Nonuniform Rational B -Spline Geometrical Representation for Aerodynamic Design of Wings," *AIAA Journal*, Vol. 39, No. 11, 2001, pp. 2033–2041. doi:10.2514/2.1206
- [2] Song, W., and Keane, A., "A Study of Shape Parameterisation Methods for Airfoil Optimization," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA, Reston, VA, 2004, pp. 2031–2038.
- [3] Painchaud-Ouellet, S., Tribes, C., Trépanier, J., and Pelletier, D., "Airfoil Shaped Optimization Using a Nonuniform Rational B -Spline Parameterization Under Thickness Constraint," *AIAA Journal*, Vol. 44, No. 10, 2006, pp. 2170–2178. doi:10.2514/1.15117
- [4] Hicks, R., and Henne, P., "Wing Design by Numerical Optimisation," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412. doi:10.2514/3.58379
- [5] Sobieczky, H., "Parametric Airfoils and Wings," *Notes on Numerical Fluid Mechanics*, Vol. 68, 1998, pp. 71–88.
- [6] Song, W., and Keane, A., "Surrogate-Based Aerodynamic Shape Optimization of a Civil Aircraft Engine Nacelle," *AIAA Journal*, Vol. 45, No. 10, 2007, pp. 2565–2574. doi:10.2514/1.30015
- [7] Robinson, G., and Keane, A., "Concise Orthogonal Representation of Supercritical Aerofoils," *Journal of Aircraft*, Vol. 38, No. 3, 2001, pp. 580–583. doi:10.2514/2.2803
- [8] Abbott, I., and Von Doenhoff, A., *Theory of Wing Sections*, Dover, New York, 1960.
- [9] Kamali, M., Ponnambalam, K., and Soulis, E., "Integration of Surrogate Optimization and PCA for Calibration of Hydrologic Models, A WATCLASS Case Study," *IEEE International Conference on Systems, Man and Cybernetics*, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 2007, pp. 2733–2737.
- [10] LeGresley, P., and Alonso, J., "Dynamic Domain Decomposition and Error Correction for Reduced Order Models," 41st AIAA Aerospace Sciences Meeting & Exhibit, Reno, NV, AIAA Paper 2003-0250, Jan. 2003.
- [11] Bui-Thanh, T., and Willcox, K., "Aerodynamic Data Reconstruction and Inverse Design Using Proper Orthogonal Decomposition," *AIAA Journal*, Vol. 42, No. 8, 2004, pp. 1505–1516. doi:10.2514/1.2159
- [12] Sacks, J., Welch, W., Mitchell, T., and Wynn, H., "Design and Analysis of Computer Experiments," *Statistical Science*, Vol. 4, No. 4, 1989, pp. 409–435. doi:10.1214/ss/1177012413
- [13] Hoyle, N., Bressloff, N., and Keane, A., "Design Optimization of a Two-Dimensional Subsonic Engine Air Intake," *AIAA Journal*, Vol. 44, No. 11, 2006, pp. 2672–2681. doi:10.2514/1.16123
- [14] Keane, A., "Statistical Improvement Criteria for Use in Multi-Objective Design Optimization," *AIAA Journal*, Vol. 44, No. 4, 2006, pp. 879–891. doi:10.2514/1.16875
- [15] Huang, D., Allen, T., Notz, W., and Miller, R., "Sequential Kriging Optimization Using Multiple-Fidelity Evaluations," *Structural and Multidisciplinary Optimization*, Vol. 32, No. 5, 2006, pp. 369–382. doi:10.1007/s00158-005-0587-0
- [16] Sakata, S., Ashida, F., and Zako, M., "Structural Optimization Using Kriging Approximation," *Computer Methods in Applied Mechanics and Engineering*, Vol. 192, Nos. 7–9, 2003, pp. 923–939. doi:10.1016/S0045-7825(02)00617-5
- [17] Forrester, A., Keane, A., and Bressloff, N., "Design and Analysis of 'Noisy' Computer Experiments," *AIAA Journal*, Vol. 44, No. 10, 2006, pp. 2331–2339. doi:10.2514/1.20068
- [18] Jones, D., "A Taxonomy of Global Optimization Methods Based on Response Surfaces," *Journal of Global Optimization*, Vol. 21, No. 4, 2001, pp. 345–383. doi:10.1023/A:1012771025575
- [19] Hollingsworth, P., and Mavris, D., "Gaussian Process Meta-Modelling: Comparison of Gaussian Process Training Methods," AIAA 3rd Annual Aviation Technology, Integration and Operations (ATIO), Denver, CO, AIAA Paper 2003-6761, Nov. 2003.
- [20] Lewin, G., and Haj-Hariri, H., "Reduced-Order Modelling of a Heaving Airfoil," *AIAA Journal*, Vol. 43, No. 2, 2005, pp. 270–283. doi:10.2514/1.8210
- [21] Siegal, S., Cohen, K., Seidel, J., and McLaughlin, T., "Proper Orthogonal Decomposition Snapshot Selection for State Estimation of Feedback Controlled Flows," 44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 2006-1400, Jan. 2006.
- [22] Hung, V., and Hien, T., "Modeling and Control of Physical Processes Using Proper Orthogonal Decomposition," *Mathematical and Computer Modelling*, Vol. 33, Nos. 1–3, 2001, pp. 223–236. doi:10.1016/S0895-7177(00)00240-5
- [23] Gross, A., and Fasel, H., "Control-Oriented Proper Orthogonal Decomposition Models for Unsteady Flows," *AIAA Journal*, Vol. 45, No. 4, 2007, pp. 814–827. doi:10.2514/1.22774
- [24] Li, G., Li, M., Azarm, S., Rambo, J., and Joshi, Y., "Optimizing Thermal Design of Data Centre Cabinets with a New Multi-Objective Genetic Algorithm," *Distributed and Parallel Databases*, Vol. 21, Nos. 2–3, 2007, pp. 167–192. doi:10.1007/s10619-007-7009-9
- [25] My-Ha, D., Lim, K., Khoo, B., and Willcox, K., "Real-Time Optimisation Using Proper Orthogonal Decomposition: Free Surface Prediction Due to Underwater Bubble Dynamics," *Computers and Fluids*, Vol. 36, No. 3, 2007, pp. 499–512. doi:10.1016/j.compfluid.2006.01.016
- [26] LeGresley, P., and Alonso, J., "Airfoil Design Optimization Using Reduced Order Models Based on Proper Orthogonal Decomposition," Fluids 2000 Conference and Exhibit, Denver, CO, AIAA Paper 2000-2545, June 2000.
- [27] Sirovich, L., "Turbulence and Dynamics of Coherent Structures Part 1: Coherent Structures," *Quarterly of Applied Mathematics*, Vol. 45, No. 3, 1987, pp. 561–571.
- [28] Lucia, D., and Beran, S., "Reduced-Order Model Development Using Proper Orthogonal Decomposition and Volterra Theory," *AIAA Journal*, Vol. 42, No. 6, 2004, pp. 1181–1190. doi:10.2514/1.10419
- [29] Jolliffe, I., *Principal Component Analysis*, Springer, New York, 2002.
- [30] Freestone, M. M., "VGK Method for Two-Dimensional Airfoil Sections," ESDU International, Data Sheet ESDU-96028, London, 1996.
- [31] Lock, R., and Williams, B., "Viscous-Inviscid Interactions in External Aerodynamics," *Progress in Aerospace Sciences*, Vol. 24, No. 2, 1987, pp. 51–171.

- doi:10.1016/0376-0421(87)90003-0
- [32] Cook, P., McDonald, M., and Firmin, M., "Aerofoil RAE-2822—Pressure Distributions, and Boundary Layer and Wake Measurements," AGARD Rept. AR-138, Neuilly-sur-Seine, France, 1979.
- [33] Barrett, T., Bressloff, N., and Keane, A., "Airfoil Shape Design and Optimization Using Multifidelity Analysis and Embedded Inverse Design," *AIAA Journal*, Vol. 44, No. 9, 2006, pp. 2051–2060. doi:10.2514/1.18766
- [34] Pound, G., and Price, A., "The Geodise OptionsMatlab Toolbox—A User's Guide, 2007, Univ. of Southampton, Southampton, England, U.K., <http://www.geodise.org/documentation/OptionsMatlab/html/index.htm> [retrieved Feb. 2010].
- [35] Söbester, A., Leary, S., and Keane, A., "On the Design of Optimization Strategies Based on Global Response Surface Approximation Models," *Journal of Global Optimization*, Vol. 33, No. 1, 2005, pp. 31–59. doi:10.1007/s10898-004-6733-1
- [36] Toal, D., Bressloff, N., and Keane, A., "Kriging Hyperparameter Tuning Strategies," *AIAA Journal*, Vol. 46, No. 5, 2008, pp. 1240–1252. doi:10.2514/1.34822
- [37] Jones, D., Schonlau, M., and Welch, W., "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, Vol. 13, No. 4, 1998, pp. 455–492. doi:10.1023/A:1008306431147
- [38] Locatelli, M., "Bayesian Algorithms for One-Dimensional Global Optimization," *Journal of Global Optimization*, Vol. 10, No. 1, 1997, pp. 57–76. doi:10.1023/A:1008294716304
- [39] Anderberg, M., *Cluster Analysis for Applications*, Academic Press, New York, 1975.

P. Beran
Associate Editor